

### Remarks

Reconsideration of the application, as previously amended, and allowance of all pending claims is respectfully requested. Claims 1-3 remain pending.

In the Office Action dated February 23, 2006, claims 1-3 were finally rejected under 35 U.S.C. 103(a) as being unpatentable over Sethuram et al. (U.S. Patent No. 5,828,903 in view of Ayanoglu et al. (U.S. Patent No. 6,122,759) in further view of Gentry et al. (U.S. Patent No. 5,778,180). Applicants respectfully, but most strenuously, traverse this rejection for the reasons cited below.

Preliminarily, and as a summary of Applicants' position, it is noted that the base patent upon which the Examiner relies to support his rejection clearly teaches the inclusion of a process step that the present inventors specifically avoid. It is through this avoidance that the data transfer process is made more efficient. Accordingly, it is seen that the patent to Sethuram et al., when properly considered, is seen to teach away from that which Applicants' have claimed as their invention.

As further evidence of the highly advanced nature of the process described and claimed in the present application, the Examiner's attention is directed to page 19, lines 17-18 wherein there is reference to real address information at the receiver end being processed via translation at the send side of the process. It is with this kind of real address information that the present process can avoid the buffering step required by Sethuram et al.

As a further help in explaining the differences between the cited base patent to Sethuram et al., it is noted that while Sethuram et al. refer to transferring data packets (cells in their language) "directly in the host memory," they do not intend by this statement that the host memory is actually an application level user's address space. The buffers to which they refer are operating system level buffers. As such, it is left to further interaction between the application level program and the operating system to arrange further transfer into the user's own storage area. There is no teaching, disclosure or suggestion that the final data destination produced by their process is something other than an area of memory that is under direct control and accessible only with the authority of the host's operating system. In

contrast, it is seem that Applicants' claims, particularly as previously amended, include a recitation that asserts that "said memory locations [are] ... application level address space locations," a clear distinction.

In this present response, the Examiner's asserted reasons for rejection are considered in the order presented in the currently presented Final Office Action.

In paragraph 2 of the Final Office Action the Examiner asserts that Applicants are attempting to distinguish their claims from the invention of Sethuram et al. via the assertion that Sethuram et al. teach that their cells (best understood as packets of information) contain header information used for subsequent processing. The Examiner effectively asserts that Applicants' claims do not preclude the inclusion of a header in their packets. However, it is a well settled principle of patent law that claims are supposed to positively set forth the collection of related elements that provide the useful device or, in this case, process. Applicants' are not required, nor is it logical that their claims set forth things that are not present in their invention.

For example, if a cited patent includes a claim for a process with steps A, B, C and D and if an applicant submits a claim for a process that produces the same result with steps A, B and C, there is no requirement in patent law for patent applicants to specifically exclude step D from their claim. It is quite possible that the data processed in the currently submitted claims does include data that is interpreted as being header or header-like information albeit at a higher level of processing. This would be a level above that associated with the transfer of the data packet. Accordingly, from the perspective of well settled patent law, it is seen that it is inappropriate for the Examiner to ask the present Applicant's to modify their claims so that the data does not include header information. Furthermore, it is a request to narrow the claims to a point which is narrower than is required. In particular, if the claims were amended to preclude the inclusion of header information, that would also prohibit any subsequent level of processing from handling information which it would interpret as being a header. It is thus seen that such a requirement is not only improper but does not fit the breadth of the claimed invention.

The Examiner makes much of Applicants' reference to Column 1, lines 33-36 of the patent to Sethuram et al. as being in their "Background Description" of the prior art. However, the same limitations are set forth throughout the cited base patent. For example, consider the following language from Column 1 of the cited patent which is titled "BRIEF SUMMARY OF THE INVENTION":

"In order to accomplish these and other objectives a system is proposed for segmenting and reassembling cells directly in the host memory. The system uses an adapter coupled between a host device, such as a computer system, and a network. The adapter includes an adapter local memory and a direct memory access (DMA) engine coupled to the host memory and performs DMA transfers between the adapter and host memory.

Virtual registers (VRs) are established in the adapter local memory to identify the location of buffers in the host memory used to store data transferred to and from the network. The DMA engine references the VRs in the adapter local memory to determine the host memory locations to access to perform the segmentation and reassembly of cells (SAR). Therefore, when a cell is received over the network, the virtual circuit is identified from the cell header information, the virtual register corresponding to the virtual circuit is accessed and the data is written by the DMA controller directly into the host memory. **When data is to be transmitted, the host device writes the data into the host memory in the buffer identified for a specified virtual circuit and the adapter is notified of the data to be transmitted.** The adapter responds by directly accessing the buffer in the host memory to extract the data for a cell, adding cell header information corresponding to the identified virtual circuit and transmitting the cell out on the network."

It is clear from the above that Sethuram et al. contemplate a process in which data that is to be transmitted is first processed in the host. Consider the process that Sethuram et al. employ for transmitting information. Clearly, their process requires that ***the data to be transmitted is first written into the host memory in the buffer identified for a specified virtual circuit*** (this being an almost exact paraphrasing of the language from the cited patent, and not simply language describing background art).

This is in direct contrast to the claimed invention in which transfer is not from some buffer but rather from desired application level address space locations. In fact, the subject recitation in Applicants' claims reads as follows: "transferring, from said second data processing system to said adapter, real address information indicating target memory

locations for said message, said memory locations being application level address space locations.” Thus, the present Applicants teach that the transfer is not to a (staging) buffer, but rather to application level address space locations identified by real (as opposed to virtual) address information.

This distinction between the claimed invention is also supported by the following passage from the cited base patent to Sethuram et al. The prefatory language used below leaves no doubt whatsoever but that the teachings are those of Sethuram et al. and not merely background art. Sethuram et al. unequivocally adopt these words as their own (see their col. 4, lines 12-25):

To transmit data, the host device 200 writes the data to be transmitted into the buffers in the host memory 208. The adapter 202 transmits the data in these buffers as 53 byte cells by retrieving small portions or segments of data directly from the host memory 208. Thus, in one embodiment, the transmit process involves the following steps: (1) the host device 200 prepares the buffers for transmission by writing data into the buffers, (2) the host device 200 informs the adapter 202 of the existence of these buffers, (3) the adapter 202 retrieves the data from the buffers in 48 byte segments and adds the five bytes of header information to form a 53 byte cell, (4) the adapter 202 transmits the cells to the network, and (5) the adapter 202 notifies the host device 200 of transmission completion. [Emphasis added herein.]

There is no doubt but that in the process of Sethuram et al., **assembly of cells** occurs prior to transmission: “...the host device 200 writes the data to be transmitted into the buffers in the host memory 208.” In Applicant’s claimed process it is precisely this step that is avoided. In Applicants’ claimed process, the data to be transmitted requires no “assembly of cells.” **Thus, Applicant’s claimed process avoids a step that is specifically taught as being required by Sethuram et al.**

However, it is still clear from the Examiner’s comments that there remains a fundamental disagreement with respect to the meaning and import of the patent to Sethuram et al. It is hoped that the present parsing of the passages found within this patent has helped to better characterize the teachings of Sethuram et al. and to set those teachings in a better light for comparison with Applicants’ claims. To whatever extent necessary, Applicants’

attorney is willing to work with the Examiner in hopes of arriving at language which both would find acceptable in view of the cited art and the breadth of the invention herein.

In further support of Applicants' position, attention is directed to column 4, lines 45-46 of the cited patent which is found under their heading for "DETAILED DESCRIPTION":

"In one embodiment, a separate buffer is allocated for all virtual circuits which have not yet been established."

This passage from Sethuram et al. clearly indicates that a receiving host system includes a single buffer for cells arriving from a plurality of sources as specified at a later time via a virtual circuit identified by a virtual register (VR) in their adapter. This passage provides even further evidence that Sethuram et al. do not contemplate a process in which arriving information packets bypass operating system level buffers. Future arriving packets that have been assigned to a virtual circuit get sent to a single catch-all buffer. This is not the targeted packet delivery system contemplated by the present inventors. In fact, it is quite the opposite: the target is linked to an uncontrollable plurality of potentially distinct sources, none of which has been "cleared" to transmit data directly into its final destination in an application level user's address space.

Those of ordinary skill in the art would not be led to a process which includes a step which is necessary in the method of Sethuram et al. to accomplish what Applicants have done without this step. The fundamental patent upon which the Examiner relies to support the rejection is actually thus seen to teach away from that which Applicants have claimed.

The other two patents cited in support of the rejection cannot and do not compensate for this fundamental difference. It is a well settled principle of patent law that art which teaches away from that which applicants claim cannot support a rejection under 35 USC § 103.

Attention is now directed to the Examiner's comments as presented in his paragraph #3 in the aforementioned Office Action. The Examiner asserts that the recitation of writing data "directly into a user's address space" does not appear in the rejected claims. It is noted that relevant claim language in the rejected claims does, however, read as follows:

“transferring, from said second data processing system to said adapter, real address information indicating target memory locations for said message, said memory locations being application level address space locations

transferring said message, from said temporary memory in said adapter, directly into said target memory locations in the memory of said second data processing system, said transfer occurring via direct memory access”

Applicants are hard pressed to understand how the Examiner can assert that “application level address space locations” are not the same as saying “directly into a user’s address space”? In this regard it is noted that there is no requirement that the language found in an applicant’s claims match the exact wording (*in haec verba* is the term of art applied) found in an applicant’s specification or elsewhere. Accordingly, Applicants vigorously disagree with the Examiner’s position as set forth in paragraph #3.

Attention is now directed to the Examiner’s comments as presented in his paragraph #4 in the aforementioned Office Action. With respect to the patent to Gentry et al. the following passage is repeated from the prior response (col. 2, lines 46-63):

“More specifically, each ATM connection may have a private pool of buffers, into which only packets for that connection will be placed. Since the pool of buffers is private, a program can be given access to its own pool. No data copying will be required for packets received into the private pool. Therefore, a packet may be directly sent to its final destination by DMA. Additionally, protected buffer descriptors prevent corruption of data with the private buffers dedicated to the data’s final destination. When a packet arrives, if there are no private buffers available, the router falls back to a common pool of buffers which are not available to the programs and thus, must be copied. Since not all connections will be able to use private buffer pools due to lack of resources, a change in the connection from the common pool of buffers to the private pool of buffers and vice versa is available. This change affects a connection while it operates. The change takes effect on the next packet to arrive.” [Emphasis added herein.]

The Examiner concludes that for “the program to be given access to its own pool, that pool must be in the program’s address space.” This is not at all true. The passage above clearly refers to private buffers. Gentry asserts that “**a program can be given access to its own [private buffer] pool.**” It is unequivocally clear from this passage that the program does not already have such control. If it had control, it would not need to be **given** control. On the

other hand, if the pool were within the address space of the program, then the program would automatically have this control and would not need to have been given it!

It is noted that, while it appears that the goals of Gentry et al. are similar to those found in the present invention, the methods employed to achieve them are different. Again, there is no writing of data directly into a user's address space. Instead, application programs are given access to separate "private buffers." **While this is a laudable goal with similar objectives, it is not the same process. There is no teaching, disclosure or suggestion that these private buffers are within the user's address space where transferred data can be used immediately without having to communicate with the operating system to assure access to data within these private buffers.** As is clearly stated in the discussion above, programs can be given access to these private buffers. In stark contrast, in the claimed invention user level applications (programs) automatically have access to the incoming data since it is placed directly into the user's address space.

With respect to the assertion that access to a private buffer pool is the same as access to one's own address space, it would appear that it would be useful to describe the address space concept more fully to better characterize that which is performable by an end user directly and that which is performed by an end user through calls to operating system level commands. The present invention does employ commands directed to the operating system level, but once these commands have been issued, data transfer occurs directly into the user's address space for immediate use without any further reference to the operating system or reference to data storage areas under control of the operating system. In the claimed invention, there is no question but that the end user already has this access because an end user's address space is accessible to the user without any further need to communicate with the operating system. **It (an application level address space location) is not an area to which the user can be given access, it is an area to which the user already has access.**

In paragraph #4, the Examiner also asserts that the claims do not refer to a user's address space. However, as pointed out above, the language "application level address space locations" is precisely the same thing. The Examiner also asserts that the claims do not preclude the use of a private pool of buffers. That is correct but neither do they have to since

Applicants herein have demonstrated from the discussion above that a private pool of buffers still refers to an area of storage that is under control of the operating system ( as per the “can be” language discussed above). The use of such private pool buffers by an application program is not precluded by the scope of the present invention. The language used in the claims herein is completely consistent with the use of the word “comprising” as not precluding the presence of other steps beyond those specifically practiced by the claimed invention.

With respect to the Examiner’s comments in paragraph #5, Applicants attorney wishes to express his willingness to discuss any future language modifications that would be potentially acceptable. It is hoped that the explanations and distinctions set forth herein fully meet the Examiner’s concerns in this area but if not, the Examiner is invited to contact the undersigned to explore areas of possible agreement.

It is also appropriate here to comment upon the cited patent to Ayanoglu et al. It is noted that, while they appear to teach the use of an acknowledgement step, **they fail to teach anything with respect to the transfer of data directly into an application level address space**. This is a feature that is also lacking in the patents to Gentry and to Sethuram et al., as pointed out above. Ayanoglu et al. are focused on the ATM (Asynchronous Transfer Method) protocol. Nowhere in that protocol is there any reference to the transfer of data directly into an application level address space. Furthermore, it is noted that there is no mention whatsoever in the patent to Ayanoglu et al. of the term “address space” or “direct memory access (DMA).

With respect to paragraphs #6 and #7, no comments are needed other than to say that the rejection is respectfully traversed.

With respect to paragraphs #8, it is most relevant to comment upon the Examiner’s comparison of Applicants’ claims with the language found in column 6, lines 22 and column 6 lines 44-46 in the patent to Sethuram et al. which are set forth respectively below for the convenience of the Examiner:



“In the present embodiment, the DMA engine 205 includes physical register block 440 which is accessed by the DMA engine to perform DMA transfers between the adapter and the buffers located in the host memory....

Initiate the DMA transfer, wherein the data in the incoming cells is directly transferred to the appropriate buffer in the host memory identified by the virtual register.” [Emphasis added herein.]

It is Applicant’s position that reference to buffers in the host memory, as taught and contemplated by Sethuram et al., is not a reference to areas of host memory within an end user’s address space.

With respect to paragraph #9, it is noted that claim 2 is a dependent claim incorporating all of the recitations of the base claim.


With respect to paragraph #10, it is noted that none of the cited documents teach, disclose or even suggest the transfer of data directly into an end user’s address space. Art that does not teach this concept cannot be used as a basis for rejecting an applicant’s claims which specifically recite it. In support of this position, Applicants specifically refer to the following language found in their claims: “**application level address space locations**.” Accordingly, it is seen that the rejection under 35 USC § 103 is not supported by the cited art.

Since it seen from the above that the indicated rejection is not supported by a proper reading of the teachings from Sethuram et al. and/or the other two cited patents, it is respectfully requested that the rejection of Applicants’ claims 1-3 based on the patent to Sethuram et al. and the other cited patents be withdrawn.

As a disclaimer it is noted that all passages provided herein from the cited art have been reproduced from data found on the Patent and Trademark Office web site. No representation is made that these inclusions have been checked against the actually language found in the patent as issued. However, Applicants’ representative has no reason to believe that they are inaccurate.

Should the Examiner wish to discuss this case with Applicants' attorney, please contact Applicants' attorney at the below listed number.

Respectfully submitted,



---

Lawrence D. Cutter  
Attorney for Applicants  
Registration No.: 28,501

Dated: April 12, 2006.

HESLIN ROTHENBERG FARLEY & MESITI P.C.  
5 Columbia Circle  
Albany, New York 12203-5160  
Telephone: (518) 452-5600  
Facsimile: (518) 452-5579